

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 March 2003 (20.03.2003)


PCT

(10) International Publication Number
WO 03/023640 A2

- (51) International Patent Classification⁷: G06F 15/20, 9/312
- (21) International Application Number: PCT/US02/28417
- (22) International Filing Date: 6 September 2002 (06.09.2002)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 60/317,860 7 September 2001 (07.09.2001) US
- (71) Applicant (for all designated States except US): SANRAD [IL/IL]; 32 Habarzel Street, 69710 Tel-Aviv (IL).
- (71) Applicant (for MW only): KATTEN, MUCHIN, ZAVIS, ROSENMAN [US/US]; 575 Madison Avenue, New York, NY 10022-2585 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): AMIR, Shai [IL/IL]; 12 Hahermon Street, 43732 Raanana (IL). KLEIN, Yaron [IL/IL]; 40 Moshe Dayan Street, 67653 Tel-Aviv (IL).
- (74) Agents: HELFGOTT, Samson et al.; Katten Muchin Zavis Rosenman, 575 Madison Avenue, New York, NY 10022-2585 (US).
- (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, SK,

[Continued on next page]

(54) Title: LOAD BALANCING METHOD FOR EXCHANGING DATA BETWEEN MULTIPLE HOSTS AND STORAGE ENTITIES, IN IP BASED STORAGE AREA NETWORK



WO 03/023640 A2

(57) Abstract: The present invention achieves load balancing of activities on storage entities, operating in a storage area network (SAN) by assignment of address to each storage entity. The storage traffic is monitored and statistics are accumulated over an interval and then used to obtain the traffic pattern of each host-storage entity link. A statistical analysis is used to determine the optimal routing map from hosts to storage entity. Other heuristics including genetic algorithms may also be used. This mapping is set on the network by assigning the entities' addresses to the storage gateways, thus does not impose any process on the hosts or the disks.



TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

- *without international search report and to be republished upon receipt of that report*

WO 03/023640

PCT/US02/28417

**LOAD BALANCING METHOD FOR EXCHANGING DATA BETWEEN
MULTIPLE HOSTS AND STORAGE ENTITIES, IN IP BASED STORAGE
AREA NETWORK**

RELATED APPLICATIONS

- 5 This application claims the benefit under 35 U.S.C 119(e) from U.S. Provisional Application No. 60/317,860 filed September 7, 2001.

BACKGROUND OF THE INVENTION

Field of Invention

The invention relates to the load balancing of traffic in an IP based Storage Area Network (SAN). More particularly, the present invention relates to dynamic assignment of IP addresses to storage end points according to network traffic.

Description of the Related Art

The Small Computer Systems Interface ("SCSI") is a popular family of protocols for communicating with I/O devices, in particular storage devices. More to the point, SCSI is the basic protocol for I/O operations between computers and storage subsystems. The SCSI protocol is, in essence, a point-to-point protocol.

Another popular protocol is the Transmission Control Protocol/Internet Protocol ("TCP/IP"). TCP/IP is the basic set of communication protocols for general data exchange between computers connected on a communication network. This is a common protocol used to communicate via the Internet.

Currently there is a convergence between the two protocols, that is, SCSI and TCP/IP. In particular, computers that communicate with their subsystems via SCSI are now tending to be interconnected via the Internet and utilizing TCP/IP to communicate with each other.

In view of this convergence, a standard (ietf draft) has been proposed for implementing SCSI over TCP. The currently proposed standard "draft-ietf-ips-iSCSI-13.txt" is available at <http://www.ietf.org/internet-drafts/draft-ietf-ips-iscsi-13.txt> [Satran. et. al IPS Internet Draft -ietf-ips-iscsi-15.txt [online], [retrieved on Sept. 5, 2002]. Retrieved from the Internet <URL:<http://www.ietf.org/internet-drafts/draft-ietf-ips-iscsi-15.txt>> (herein expressly incorporated by reference).

The ietf draft aims to be fully compliant with the requirements laid out in the SCSI Architecture Model-2 Document (Project Leader Charles Monia , Technical Editor John P. Scheible Working Draft [online] [retrieved on Sept. 5, 2002]. Retrieved from the Internet <<ftp://ftp.t10.org/t10/drafts/s2p/s2p-r07b.pdf>> also herein expressly incorporated by reference). This SCSI standard provides for transmitting SCSI commands between SCSI devices, over TCP/IP connections. Conventionally, SCSI devices that communicate with each other must be connected via the same SCSI bus. The proposed ietf draft permits SCSI devices, which are not connected via the same SCSI bus to communicate with each other via the Internet, utilizing the TCP/IP protocol. According to the ietf draft, SCSI devices that are separated from one another even by great distances may communicate with each other. The ietf draft seeks to describe an implementation of SCSI that operates on top of TCP as the transport protocol.

The following briefly defines some SCSI conceptual terms. The end point of a typical SCSI command is a "logical unit" (LUN). Logical units include, for example, hard drives, tape drives, CD and DVD drives, printers and processors. A collection of logical units is referred to as a "target" and is directly addressable on the network. In a client-server model, the target corresponds to the server. An "initiator" creates and sends SCSI commands to the target. In the client-server model, the initiator corresponds to the client.

A typical SCSI command results in a command phase, data phase and a response phase. In the data phase, information travels either from the initiator to the target (for example, a WRITE command), or from the target to the initiator (for example, a READ command). In the response phase, the target returns the final status of the operation, including any errors. A response signal, is the end of a typical SCSI command.

The ietf draft defines a method of encapsulating those SCSI commands, SCSI Data and SCSI Responses over TCP/IP. It also defines methods to ensure the reliability of the transfer from the initiator to the target. This includes commands' numbering, commands' tagging and acknowledgment. Also it provides means of securing data and authenticating initiators and targets to each other.

A storage gateway is a device that connects IP networks and Storage Networks (e.g., parallel SCSI or Fibre Channel). On one side it connects to initiators via an IP network and on the other side it connects to targets comprising storage devices via for example, the storage bus. It receives encapsulated SCSI Commands and/or data

WO 03/023640

PCT/US02/28417

from the initiators, decapsulates them and transfers the raw SCSI command/data to the storage device. After completion of the command, it sends back to the initiator a status signal and/or data over the IP network.

A storage gateway can include more features than simple connectivity. Those
5 can include Virtualization, disk concatenation, striping, mirroring and so on. The exposed LUNs can be any function of the physical LUNs. The present invention is independent of those features.

BRIEF SUMMARY OF THE INVENTION

10 It is the intent of this invention to provide a method, system and/or means for balancing traffic on a plurality of storage entities in a storage area network. Some examples of storage entities include physical disks, virtual disks, clusters of physical disks, clusters of virtual disks and/or combinations of virtual and physical disks.
Examples of a storage area network include an IP based network or a MAC layer based
15 network.

FIG. 1 is an illustration of a SAN over IP. The host servers (10) in the upper part are the initiators in the client-server model. Any number of host servers may be included, for example, as shown in FIG. 1 Host 1 to Host n. The storage endpoints (30) in the lower part are the target LUNs. FIG. 1 shows the LUNs as Storage 1 to Storage m
20 as an example. However the LUNs may include, for example, hard drives, tape drives, CD and DVD drives, printers and processors.

An intermediate layer is shown in FIG. 1 and may include a combination of, for example, Gateways (20). As illustrated in FIG. 1 any combination of Gateways (20) may be included. These are shown for convenience as Gateway 1 to Gateway k. The
25 topology in this SAN over IP of FIG. 1 has any-to-any connectivity in the storage layer (35). The any-to-any connectivity can connect to any storage device.

According to an embodiment of the invention, the mapping is set in the network by assigning an IP address of each storage entity. The IP address may be unique for each storage entity.

30 According to another embodiment of the invention, each gateway listens on the IP addresses of the storage entities it is responsible for. There is no intersection of IP addresses between the gateways and each storage entity has one gateway that is responsible for it. Thus a one-to-one routing map is implemented.

According to another embodiment of the invention, a statistical server accumulates and processes statistics of data traffic in the network from each host – storage entity route over a time interval.

According to another embodiment of the invention, a statistical server is calculating the optimal routing map of each host – storage entity. The optimal routing map is the solution to an optimization problem, where each solution (mapping) results a cost function combined from the traffic variance (balance) and minimal changes from the previous mapping.

According to another embodiment of the invention, the new routing map is set to the gateways. The server sets to each gateway the IP addresses it should listen via.

The present invention proposes to take advantage of a unique addressing mechanism, to simplify the logic to a great extent reducing bottlenecks in the traffic and to balance traffic loads in the network. Thus, provide faster access and higher throughput from servers to the storage devices. Examples of a unique addressing mechanism may include IP, MAC but is not limited to these schemes.

BRIEF DESCRIPTION OF THE DRAWINGS

The above-mentioned and other advantages and features of the present invention will become more readily apparent from the following detailed description and the accompanying drawings, in which:

Fig 1 is a scheme of a system topology. The scheme describes an IP network (25) with n host computers (10) and k gateways (20). The gateways (20) are connected with an any-to-any connection (35) to m storage entities (30);

Fig 2 is an example of a particular routing. It shows the connectivity (240) from Host 2 (210) to storage entity 3 (235) via gateway 1 (220);

Fig 3 is a scheme of a system topology with the addition of the statistics' server (360). The scheme describes an IP network (325) that includes n host computers (310) and k gateways (320). The gateways (320) are connected with an any-to-any connection to m storage entities (330);

Fig 4 is a scheme of another topology. The topology includes an IP network (425) connected with n host computers (410) and k gateways (420). The gateways

WO 03/023640

PCT/US02/28417

(420) are connected (455) to the storage entities (468) in a Fibre Channel network (465) routed by FC switch or FC hub (450); and

Fig 5 is a scheme of a topology containing n host computers (510), k proxies (520) and m storage endpoints with IP interface (570).

5

DETAILED DESCRIPTION OF THE INVENTION

The following detailed description of the preferred embodiment includes many specific details. The inclusion of such details is for the purpose of illustration only and should not be understood to limit the invention.

10 FIG. 1 is an embodiment of n host servers (initiators) (10) connected to k Gateways (20) via IP network (25). The k gateways (20) are connected to m storage devices (30) with an any-to-any connectivity (35). Each host can reach each disk via each one of the gateways (20).

15 For purposes of illustration for a particular embodiment, the TCP/IP backbone including routers could be part of the global Internet, an Intranet, or any network utilizing TCP/IP. It should be noted that the invention can be applied to any system utilizing SCSI buses which are interconnected via a TCP/IP protocol. Thus, the invention could easily be implemented in a variety of network architectures.

20 In this example a configuration will be used in the following discussion of the method and system according to the invention. Reference is made to FIG. 1, in connection with describing the behavior of an example network.

25 The goal of the invention is to create a routine mapping from each host server (10) to each disk (30). By creating this mapping the gateway (20) is set through which each host computer (10) connects to each disk (30), such that the traffic in the gateways (20) on the storage sub-network side (20) (35) (30) is balanced.

A mapping of each host computer (10) to each disk (30) is of the form of a matrix:

$$\begin{aligned}
 & \text{map} \left(\begin{matrix} \text{disk}_1 & \text{disk}_2 & \wedge & \text{disk}_m \end{matrix} \right), \left(\begin{matrix} \text{host}_1 \\ \text{host}_2 \\ M \\ \text{host}_n \end{matrix} \right) = \\
 & \left(\begin{matrix} \text{Gateway}_{1,1} & \text{Gateway}_{1,2} & \wedge & \text{Gateway}_{1,m} \\ \text{Gateway}_{2,1} & \text{Gateway}_{2,2} & \wedge & \text{Gateway}_{2,m} \\ M & M & O & M \\ \text{Gateway}_{n,1} & \text{Gateway}_{n,2} & \wedge & \text{Gateway}_{n,m} \end{matrix} \right)
 \end{aligned}$$

where

$$\text{Gateway}_{i,j} \in \{\text{gateway}_1, \dots, \text{gateway}_k\}$$

5 is a gateway that connects from host_i to disk_j .

The traffic i,j from each host computer to a disk is measured in a time interval, is assembled. This assembling may be performed by a specific statistic server (350) as shown in FIG. 3 or a priority and entered into a matrix of the form:

$$10 \quad \text{traffic} \left(\begin{pmatrix} \text{disk}_1 & \text{disk}_2 & \wedge & \text{disk}_m \end{pmatrix}, \begin{pmatrix} \text{host}_1 \\ \text{host}_2 \\ \vdots \\ \text{host}_n \end{pmatrix} \right) =$$

$$\begin{pmatrix} \text{traffic}_{1,1} & \text{traffic}_{1,2} & \wedge & \text{traffic}_{1,m} \\ \text{traffic}_{2,1} & \text{traffic}_{2,2} & \wedge & \text{traffic}_{2,m} \\ \vdots & \vdots & \text{O} & \vdots \\ \text{traffic}_{n,1} & \text{traffic}_{n,2} & \wedge & \text{traffic}_{n,m} \end{pmatrix}$$

Where $\text{traffic}_{i,j}$ is the amount of data transferred between host_i and disk_j in the time interval.

$$15 \quad \text{Gtraffic} = \left(\begin{pmatrix} \text{disk}_1 & \text{disk}_2 & \wedge & \text{disk}_m \end{pmatrix}, \begin{pmatrix} \text{Gateway}_1 \\ \text{Gateway}_2 \\ \vdots \\ \text{Gateway}_k \end{pmatrix} \right) =$$

$$\begin{pmatrix} \text{Gtraffic}_{1,1} & \text{Gtraffic}_{1,2} & \wedge & \text{Gtraffic}_{1,m} \\ \text{Gtraffic}_{2,1} & \text{Gtraffic}_{2,2} & \wedge & \text{Gtraffic}_{2,m} \\ \vdots & \vdots & \text{O} & \vdots \\ \text{Gtraffic}_{k,1} & \text{Gtraffic}_{k,2} & \wedge & \text{Gtraffic}_{k,m} \end{pmatrix}$$

where $\text{Gtraffic}_{i,j}$ is the traffic from gateway_i and disk_j such that:

$$20 \quad \text{Gtraffic}_{i,j} = \text{Gtraffic}(\text{Gateway}_i, \text{disk}_j) = \sum_{\text{Gateway}_i, j = \text{gateway}_i} \text{traffic}_{i,j}$$

k is:

$$\text{GatewayTraffic} = \begin{pmatrix} \text{Gtraffic}_1 \\ \text{Gtraffic}_2 \\ \vdots \\ \text{Gtraffic}_k \end{pmatrix} = \begin{pmatrix} \sum_{i=1}^m \text{Gtraffic}_{1,i} \\ \sum_{i=1}^m \text{Gtraffic}_{2,i} \\ \vdots \\ \sum_{i=1}^m \text{Gtraffic}_{k,i} \end{pmatrix}$$

WO 03/023640

PCT/US02/28417

The optimal mapping is a map matrix that will minimize the cost function:

$$\min(w_{\text{var}} \cdot \text{var}(\text{GatewayTraffic}) + w_{\text{change}} \cdot \text{change})$$

where *var* is the variance, *change* is the number of changes from the previous mapping and w_{var} and w_{change} are the weights assigned to the variance and changes respectively.

The solution for the optimization problem assigns each route from host i (10) to storage device j (30) $\text{Gateway}_{i,j}$ to a physical gateway l (20), i.e

$$\text{Gateway}_{i,j} \rightarrow \text{gateway}_l, 1 \leq i \leq m, 1 \leq l \leq k$$

The possible mappings combinations are k^{m^n} , thus the complexity of the optimization problem is $O(\text{map}) = k^{m^n}$.

where k is the number of gateways (20), m is the number of disks (30) and n is the number of host computers (10).

In networks, where $O(\text{map})$ is computationally feasible, all alternatives are checked (cost function is calculated) and the best mapping is selected. The feasibility may be related to the size of the network or the availability of computational power, for example. In addition to computing all alternatives, heuristic algorithms can be applied such as simulated annealing, greedy algorithms, economic optimizations, genetic algorithms or tabular search.

If a Storage Name Server (SNS) resides in the network, it contains the mapping. Each host computer (10) queries the IP Network periodically for the gateway (20) that connects it to the desired disk (30).

Each gateway (20) listens on the corresponding IP addresses of the routes it is assigned to. That is, if host (10) i ($1 \leq i \leq n$) is routed to storage entity (30) j ($1 \leq j \leq m$) via gateway (20) l ($1 \leq l \leq k$), gateway l listens on the IP address of storage entity j . Thus the mapping is implemented transparently for the host computers (10) and the disk (30).

In an embodiment of this method and system for balancing traffic on a plurality of host storage entities with selectively optimized routing comprising physical disks, virtual disks, clusters of physical disks, clusters of virtual disks as well as combinations of virtual and physical disks in an IP based storage area network includes constructing a network that has at least one host of several hosts, at least one of several of a storage entities, and at least one of several storage gateways;

In the embodiment a unique gateway is selected for connecting between the host entity to the storage entity. Routing occurs via a host-gateway-storage entity. Further, assigning an unique IP address to each of at least one of several storage entities such that said host recognizes a storage entity by IP address. Then accumulating at least one of several of a statistical reading of data traffic in said network from each host-gateway-storage entity route over a time interval by a statistical server. The statistical server may process the statistical reading. This may be utilized in calculating an optimal routing map of each of the host gateway-storage-entity.

The optimal routing map includes a an optimization function where each of at least one of several mapping solutions results from a cost function combined from a traffic variance balance and a reading of minimal changes for an a priori mapping;

Further determining at least one of several of a mapping setting in the network by assigning the unique IP address of each host-gateway-storage entity to a storage gateway with each of one of several gateways listening for new IP addresses set to each gateway; and implementing an optimized routing mapping.

Also in another embodiment the accumulating of statistics is done a-priory by at least a network manager.

In a further embodiment the IP addresses assignment is implemented with a Storage Name Server configuration. Also the assignments may be determined by a host querying the server periodically for the IP address of the storage entity and granting the IP address according to the routing map.

In a further embodiment the routing map is implemented by the storage gateways with redirections of the storage transaction commands.

In a further embodiment the network is constructed from at least one of several hosts, several proxies and several IP storage endpoints comprised of at least one of a storage entity with IP interface or a storage gateway.

The optimal routing may also comprise an optimization determined by a fitness function including at least one of a genetic algorithm; a simulated annealing or a tabular search.

Fig 4 shows an embodiment of another topology. An IP network (425) is connected with n host computers (410) and k gateways (420). The gateways (420) are connected (455) to the storage entities (468) in a Fibre Channel network (465) routed by FC switch or FC hub (450). The storage entities (468) may include, for example, physical disks, virtual disks, clusters of physical disks, clusters of virtual disks as

WO 03/023640

PCT/US02/28417

well as combinations of virtual and physical disks. The system implements selectively optimized routing.

This system is constructed to include at least one of several gateways (420); at least one of several storage databases; an IP network (425); at least one of several statistical servers (450); at least one of several FC switches or hubs (450) in any combination. Also included as shown in FIG. 5 are k proxies (520), thus the configuration may also include at least one of several proxies; and an interconnection optimally routed to at least one of several storage databases (468 FIG. 4). The storage databases may further include IP enabled storage databases (570 FIG. 5).

In another embodiment a system for balancing traffic on a plurality of host storage entities with selectively optimized routing in an IP based storage area network also includes any combination of physical disks, virtual disks, clusters of physical disks, clusters of virtual disks as well as combinations of virtual and physical disks.

This system further includes at least one of several hosts linking to an IP network, a statistical service, at least one of several gateways, interconnecting at least one of at least one FC switch an interconnection to optimally routed to at least one served of several storage databases.

In this embodiment the storage databases may further include IP enabled storage databases; an IP network; at least one of several statistical servers; at least one of several FC switches; at least one of several FC hubs; at least one of several proxies; and an interconnection optimally routed to at least one of several storage databases; said storage databases further comprising IP enabled storage databases.

While this invention has been described in conjunction with the specific embodiments outlined above, many alternatives, modifications and variations will be apparent to those skilled in the art. Accordingly, the preferred embodiments of the invention set forth above are intended to be illustrative, and not limiting. Various changes may be made without parting from the spirit and the scope of the invention as defined in the following claims.

What is claimed is:

1. A method for balancing traffic on a plurality of host storage entities with selectively optimized routing in an storage area network, including a network of at least one host and at least one storage gateway, said method comprising the steps of:
 - 5 selecting a unique gateway for connecting between a host entity to a storage entity;
 - routing via a host - gateway - storage entity;
 - assigning an IP address to at least one of several storage entities such that said host recognizes a storage entity by IP address;
 - 10 accumulating at least one of several of a statistical readings of data traffic in said network from each host-gateway-storage entity route over a time interval by a statistical server;
 - processing said statistical reading;
 - calculating an optimal routing map of each of said host gateway-storage-entity
 - 15 said optimal routing map including a function of an optimization function where each of at least one of several mapping solutions results from a cost function combined from a traffic variance balance and a reading of minimal changes for an apriori mapping; determining at least one of several of a mapping setting in said network by assigning said IP address of each host-gateway-storage entity to a storage gateway with each of
 - 20 one of several gateways listening for new IP addresses set to each gateway; and implementing an optimized routing mapping.
2. A method as recited in claim 1, wherein the assigning step assigns a unique IP address.
 3. A method as recited in claim 1, wherein the processing step is performed by a
25 statistical server.
 4. A method as recited in claim 1, wherein the step of accumulating of statistics in the method of claim 1 is done a-priory by at least a network manager.
 5. A method as recited in claim 1, wherein the IP addresses assignment is implemented with a Storage Name Server configuration therein further said assignments
30 is determined by the step of :
 - a host querying the server periodically for the IP address of the storage entity and granting the IP address according to the routing map.

WO 03/023640

PCT/US02/28417

6. A method as recited in claim 1, wherein the routing map is implemented by the storage gateways with redirections of the storage transaction commands.
7. A method as recited in claim 1, wherein the network is constructed from at least one of several hosts, several proxies and several IP storage endpoints including at least one of a storage entity with IP interface or a storage gateway.
 - 5
8. A method as recited in claim 1, wherein said optimal routing further includes an optimization determined by a fitness function having at least one of a genetic algorithm; a stochastic simulation; a simulated annealing; a tabular search; an economic optimization, wherein the economic optimization further includes at least one of a cost minimization; a substitution based on a selection of a cluster of at least one of several of a host; at least one of several of a gateway; and at least one of several of a storage entity.
 - 10
9. A method as recited in claim 1, wherein the host storage entity includes at least one of a host storage entity having at least one of several of physical disks, virtual disks, cluster of physical disks, cluster of virtual disks or cluster of virtual and physical disks.
 - 15
10. A system for balancing traffic on a plurality of host storage entities with selectively optimized routing comprising physical disks, virtual disks, clusters of physical disks, clusters of virtual disks as well as combinations of virtual and physical disks in an IP based storage area network;
 - 20
 - at least one of several gateways;
 - at least one of several storage databases;
 - an IP network;
 - at least one of several statistical servers;
 - 25
 - at least one of several FC switches;
 - at least one of several FC hubs;
 - at least one of several proxies; and
 - an interconnection optimally routed to;
 - at least one of several storage databases;
 - 30
 - said storage databases further comprising IP enabled storage databases.
11. A system for balancing traffic on a plurality of host storage entities with selectively optimized routing comprising physical disks, virtual disks, clusters of physical disks, clusters of virtual disks as well as combinations of virtual and physical disks in an IP based storage area network; further consisting of

at least one of several hosts linking to an;
an IP network;
a statistical service;
at least one of several gateways interconnecting to
5 at least one of a FC switch;
an interconnection optimally routed to
at least one served of several storage databases;
said storage databases further comprising IP enabled storage databases.
an IP network;
10 at least one of several statistical servers;
at least one of several FC switches;
at least one of several FC hubs;
at least one of several proxies; and
an interconnection optimally routed to
15 at least one of several storage databases;
said storage databases further comprising IP enabled storage databases.

WO 03/023640

PCT/US02/28417

1/5




FIG. 1




FIG. 2

WO 03/023640

PCT/US02/28417

3/5




FIG. 3

4/5




FIG. 4

WO 03/023640

PCT/US02/28417

5/5




FIG. 5

This Page Blank (uspto)